

Amendments to the Claims:

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

1 1. (Previously Presented) In a data storage network, a method of maintaining
2 data coherency using two or more array management functions (AMFs) that are able to
3 concurrently access a redundancy group, the redundancy group including a plurality of resources,
4 the method comprising:

5 receiving a request from a host, by a first one of the AMFs, to perform a first
6 operation on data stored on a first one of the resources;

7 broadcasting a message from the first AMF to the other AMFs sharing access to
8 the first resource so as to acquire access to the first resource;

9 performing the first operation on the data by the first AMF; and

10 sending replication and state data from the first AMF to the other AMFs
11 concurrently with performing the operation such that if the first AMF fails while performing any
12 steps of the operation, one of the other AMFs is able to complete the operation using said sent
13 replication and state data.

1 2. (Original) The method of claim 1, wherein the first operation is a write
2 operation, the method further comprising receiving from the host the data to be written to the
3 first resource by the first AMF.

1 3. (Original) The method of claim 2, wherein the broadcast message is a
2 write invalidate request, and wherein, responsive to the invalidate request, each of the AMFs
3 sharing access to the resource invalidates corresponding data stored in its cache.

4. (Canceled)

5. (Original) The method of claim 1, wherein the first operation is a read operation.

6. (Original) The method of claim 5, further comprising performing a search of the cache of each of the AMFs sharing access to the first resource for a copy of the data requested in the read request.

7. (Original) The method of claim 6, further comprising reading the requested data from the first resource if none of the AMFs sharing access respond with a copy of the requested data.

8. (Original) The method of claim 6, further comprising receiving the requested data from one of the AMFs sharing access to the resource.

9. (Original) The method of claim 5, wherein the broadcast message identifies the data in the request, the method further comprising searching the cache of each of the AMFs sharing access to the first resource for the identified data, and forwarding the identified data to the first AMF if found in the cache.

10. (Original) The method of claim 1, further comprising determining whether the data identified in the request is shared by one or more other AMFs.

11. (Previously Presented) In a data storage network, a method of maintaining data coherency using two or more array management functions (AMFs) that are able to concurrently access a redundancy group, the redundancy group including a plurality of resources, the method comprising:

receiving a request from a host, by a first one of the AMFs, to perform a first operation on data stored on a first one of the resources;

determining from the request whether the identified data is shared by one or more
of the other AMFs; and

if so:

broadcasting a message from the first AMF to the other AMFs sharing access to
the first resource so as to acquire access to the first resource; and

performing the first operation on the data by the first AMF; and

if not:

performing the first operation on the data by the first AMF; and

in both cases, sending replication and state data from the first AMF to the other
AMFs concurrently with performing the first operation such that if the first AMF fails while
performing any steps of the first operation, one of the other AMFs is able to complete the first
operation using said sent replication and state data.

12. (Original) The method of claim 11, wherein the request is one of a write
data request and a read data request.

13. (Canceled)

14. (Original) The method of claim 11, wherein the broadcast message is a
write invalidate request, and wherein, responsive to the invalidate request, each of the AMFs
sharing access to the resource invalidates corresponding data stored in its cache.

15. (Previously Presented) A data storage network, comprising:

two or more controllers communicably coupled to a redundancy group, each
controller having one or more array management functions (AMFs), wherein two or more of said

AMFs are able to concurrently access the redundancy group, the redundancy group including a plurality of resources, each controller further including:

means for receiving a request from a host to perform an operation on data stored on a first one of the resources by a first AMF;

means for broadcasting a message from the first AMF to other AMFs sharing access to the first resource so as to acquire access to the first resource;

means for performing the first operation on the data by the first AMF; and

means for sending replication and state data from the first AMF to the other AMFs concurrently with performing the operation such that if the first AMF fails while performing any steps of the operation, one of the other AMFs is able to complete the operation using said sent replication and state data.

16. (Previously Presented) A data storage network, comprising:

two or more controllers communicably coupled to a redundancy group, each controller having one or more array management functions (AMFs), wherein two or more of said AMFs are able to concurrently access the redundancy group, the redundancy group including a plurality of resources, each controller further including:

means for receiving a request from a host to perform a first operation on data stored on a first one of the resources by a first AMF;

means for determining from the request whether the identified data is shared by one or more of the other AMFs;

means for broadcasting a message from the first AMF to the other AMFs sharing access to the first resource so as to acquire access to the first resource if the data is shared;

12 means for performing the first operation on the data by the first AMF; and
13 means for sending replication and state data from the first AMF to the other
14 AMFs concurrently with performing the operation such that if the first AMF fails while
15 performing any steps of the operation, one of the other AMFs is able to complete the operation
16 using said sent replication and state data.